

Ryan Stephens  
Ron Plew  
Arie D. Jones

**Fourth Edition**

Sams **Teach Yourself**

**SQL**®

in **24**  
**Hours**

**SAMS**



## Sams Teach Yourself SQL® in 24 Hours, Fourth Edition

Copyright © 2008 by Sams Publishing

All rights reserved. No part of this book shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

ISBN-13: 978-0-672-33018-6

ISBN-10: 0-672-33018-0

*Library of Congress Cataloging-in-Publication Data*

Stephens, Ryan K.

*Sams teach yourself SQL in 24 hours / Ryan Stephens, Ron Plew, Arie D. Jones. – 4th ed.*

*p. cm. – (Sams teach yourself in 24 hours)*

*On t.p. of earlier ed. Ronald R. Plew's name appeared first.*

*Includes indexes*

*ISBN 978-0-672-33018-6 (pbk.)*

*1. SQL (Computer program language) I. Plew, Ronald R. II. Jones, Arie. III. Plew, Ronald R. Sams teach yourself SQL in 24 hours. IV. Title.*

*QA76.73.S67P554 2008*

*005.75'6-dc22*

2008016630

Printed in the United States of America

First Printing May 2008

### Trademarks

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Sams Publishing cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

### Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an “as is” basis. The authors and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book.

### Bulk Sales

Sams Publishing offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales. For more information, please contact

**U.S. Corporate and Government Sales**

**1-800-382-3419**

**corpsales@pearsontechgroup.com**

For sales outside of the U.S., please contact

**International Sales**

**international@pearson.com**

**Associate  
Publisher**

*Mark Taub*

**Acquisitions Editor**

*Trina MacDonald*

**Development  
Editor**

*Michael Thurston*

**Managing Editor**

*Patrick Kanouse*

**Project Editor**

*Mandie Frank*

**Copy Editor**

*Heather Wilkins*

*Editorial Services*

**Indexer**

*Heather McNeil*

**Proofreader**

*Matt Purcell*

**Technical Editor**

*Steve Cvar*

**Publishing  
Coordinator**

*Olivia Basegio*

**Designer**

*Gary Adair*

**Composition**

*Bronkella Publishing*



### This Book Is Safari Enabled

The Safari® Enabled icon on the cover of your favorite technology book means the book is available through Safari Bookshelf. When you buy this book, you get free access to the online edition for 45 days.

Safari Bookshelf is an electronic reference library that lets you easily search thousands of technical books, find code samples, download chapters, and access technical information whenever and wherever you need it.

To gain 45-day Safari Enabled access to this book:

- ▶ Go to <http://www.informit.com/onlineedition>
- ▶ Complete the brief registration form
- ▶ Enter the coupon code SNRG-H3DQ-9TLV-YNQV-Q3VU

If you have difficulty registering on Safari Bookshelf or accessing the online edition, please email [customer-service@safaribooksonline.com](mailto:customer-service@safaribooksonline.com).

# Introduction

Welcome to the world of relational databases and SQL! This book is written for those self-motivated individuals out there who would like to get an edge on relational database technology by learning the Structured Query Language—SQL. This book was written primarily for those with very little or no experience with relational database management systems using SQL. This book also applies to those who have some experience with relational databases but need to learn how to navigate within the database, issue queries against the database, build database structures, manipulate data in the database, and more. This book is not geared toward individuals with significant relational database experience who have been using SQL on a regular basis.

## What This Book Intends to Accomplish

This book was written for individuals with little or no experience using SQL or those who have used a relational database, but their tasks have been very limited within the realm of SQL. Keeping this thought in mind, it should be noted up front that this book is strictly a learning mechanism, and one in which we present the material from ground zero and provide examples and exercises with which to begin to apply the material covered. This book is not a complete SQL reference and should not be relied on as a sole reference of SQL. However, this book combined with a complete SQL command reference could serve as a complete solution source to all of your SQL needs.

## What We Added to This Edition

This edition contains the same content and format as the first through third editions. We have been through the entire book, searching for the little things that could be improved to produce a better edition. We have also added concepts and commands from the new SQL standard, SQL:2003, to bring this book up to date, making it more complete and applicable to today's SQL user. The most important addition was the use of MySQL for hands-on exercises. By using an open source database such as MySQL, all readers have equal opportunity for participation in hands-on exercises.

## What You Need

You might be wondering, what do I need to make this book work for me? Theoretically, you should be able to pick up this book, study the material for the current hour, study the examples, and either write out the exercises or run them on a relational database server. However, it would be to your benefit to have access to a relational database system to which to apply the material in each lesson. The relational database to which you have access is not a major factor because SQL is the standard language for all relational databases. Some database systems that you can use include Oracle, Sybase, Informix, Microsoft SQL Server, Microsoft Access, MySQL, and dBASE.

## Conventions Used in This Book

For the most part, we have tried to keep conventions in this book as simple as possible.

Many new terms are printed in italics.

In the listings, all code that you type in (input) appears in boldface monospace. Output appears in standard monospace. Any code that is serving as a placeholder appears in *italic monospace*.

SQL code and keywords have been placed in uppercase for your convenience and general consistency. For example:

```
SELECT * FROM PRODUCTS_TBL;
```

```
PROD_ID    PROD_DESC                                COST
-----
11235     WITCHES COSTUME                          29.99
222       PLASTIC PUMPKIN 18 INCH                   7.75
13        FALSE PARAFFIN TEETH                      1.1
90        LIGHTED LANTERNS                          14.5
15        ASSORTED COSTUMES                         10
9         CANDY CORN                                1.35
6         PUMPKIN CANDY                             1.45
87        PLASTIC SPIDERS                           1.05
119       ASSORTED MASKS                             4.95
```

9 rows selected.

The following special design features enhance the text:

There are syntax boxes to draw your attention to the syntax of the commands discussed during each hour.

```
SELECT [ ALL | * | DISTINCT COLUMN1, COLUMN2 ]
FROM TABLE [ , TABLE2 ];
```

Notes are provided to expand on the material covered in each hour of the book.

***By the  
Way***

Cautions are provided to warn the reader about “disasters” that could occur and certain precautions that should be taken.

***Watch  
Out!***

Tips are also given to supplement the material covered during appropriate hours of study.

***Did you  
Know?***

## ANSI SQL and Vendor Implementations

One thing that is difficult about writing a book like this on standard SQL is that although there is an ANSI standard for SQL, each database vendor has its own implementation of SQL. With each implementation come variations from the actual standard, enhancements to the standard, and even missing elements from the standard.

The expected question is, “Because there is an ANSI standard for SQL, what is so difficult about teaching standard SQL?” The answer to this question begins with the statement that ANSI SQL is just that: a standard. ANSI SQL is not an actual language. To teach you SQL, we had to come up with examples and exercises that involve using one or more implementations of SQL. Because each vendor has its own implementation with its own specifications for the language of SQL, these variations, if not handled properly in this book, could actually cause confusion concerning the syntax of various SQL commands. Therefore, we have tried to stay as close to the ANSI standard as possible, foremost discussing the ANSI standard and then showing examples from different implementations that are very close, if not the same, as the exact syntax that ANSI prescribes.

We have, however, accompanied examples of variations among implementations with notes for reminders and tips on what to watch out for. Just remember this: Each implementation differs slightly from other implementations. The most important thing is that you understand the underlying concepts of SQL and its commands. Although slight variations do exist, SQL is basically the same across the board and is very portable from database to database, regardless of the particular implementation.

## Understanding the Examples and Exercises

We have chosen to use MySQL for most of the examples in this book due to its high compliance to the ANSI standard; however, we have also shown examples from Oracle, Sybase, Microsoft SQL Server, and dBASE.

## Sams Teach Yourself SQL in 24 Hours

The use of MySQL for hands-on exercises was chosen so that all readers may participate, with minimal confusion in converting SQL syntax into the proper syntax of the database each reader is using. MySQL was chosen for exercises because it is an open source database (free), it is easy to install, and its syntax is very similar to that of the ANSI Standard. Additionally, MySQL is compatible with most operating system platforms.

In Appendix B, “Using MySQL for Exercises,” we show you how to obtain and install MySQL. After it is installed on your computer, MySQL can be used for most of the exercises in this book. Unfortunately, because MySQL is not fully ANSI SQL compliant, MySQL exercises are not available for every subject.

As stated, some differences in the exact syntax exist among implementations of SQL. For example, if you attempt to execute some examples in this book, you might have to make minor modifications to fit the exact syntax of the implementation that you are using. We have tried to keep all the examples compliant with the standard; however, we have intentionally shown you some examples that are not exactly compliant. The basic structure for all the commands is the same. To learn SQL, you have to start with an implementation using practical examples. For hands-on practice, we use MySQL. If you have access to another database implementation such as Oracle, we encourage its use for hands-on exercises. You should be able to emulate the database and examples used in this book without much difficulty. Any adjustments that you might have to make to the examples in this book to fit your implementation exactly will only help you to better understand the syntax and features of your implementation.

Good luck!

## HOUR 3

# Managing Database Objects

In this hour, you learn about database objects: what they are, how they act, how they are stored, and how they relate to one another. Database objects are the underlying backbone of the relational database. These *objects* are logical units within the database that are used to store information and are referred to as the *back-end database*. The majority of the instruction during this hour revolves around the table, but keep in mind that there are other database objects, many of which are discussed in later hours of study.

---

### ***The highlights of this hour include:***

- ▶ An introduction to database objects
- ▶ An introduction to the schema
- ▶ An introduction to the table
- ▶ A discussion of the nature and attributes of tables
- ▶ Examples for the creation and manipulation of tables
- ▶ A discussion of table storage options
- ▶ Concepts on referential integrity and data consistency

## **What Are Database Objects?**

A *database object* is any defined object in a database that is used to store or reference data. Some examples of database objects include tables, views, clusters, sequences, indexes, and synonyms. The table is this hour's focus because it is the primary and simplest form of data storage in a relational database.

## What Is a Schema?

A *schema* is a collection of database objects (as far as this hour is concerned—tables) associated with one particular database username. This username is called the *schema owner*, or the owner of the related group of objects. You may have one or multiple schemas in a database. The user is only associated with the schema of the same name and often the terms will be used interchangeably. Basically, any user who creates an object has just created it in her own schema unless she specifically instructs it to be created in another one. So, based on a user's privileges within the database, the user has control over objects that are created, manipulated, and deleted. A schema can consist of a single table and has no limits to the number of objects that it may contain, unless restricted by a specific database implementation.

Say you have been issued a database username and password by the database administrator. Your username is USER1. Suppose you log on to the database and then create a table called EMPLOYEE\_TBL. According to the database, your table's actual name is USER1.EMPLOYEE\_TBL. The schema name for that table is USER1, which is also the owner of that table. You have just created the first table of a schema.

The good thing about schemas is that when you access a table that you own (in your own schema), you do not have to refer to the schema name. For instance, you could refer to your table as either one of the following:

```
EMPLOYEE_TBL  
USER1.EMPLOYEE_TBL
```

The first option is preferred because it requires fewer keystrokes. If another user were to query one of your tables, the user would have to specify the schema, as follows:


```
USER1.EMPLOYEE_TBL
```

In Hour 20, "Creating and Using Views and Synonyms," you learn about the distribution of permissions so that other users can access your tables. You also learn about synonyms, which allow you to give a table another name so you do not have to specify the schema name when accessing a table. Figure 3.1 illustrates two schemas in a relational database.



There are, in Figure 3.1, two user accounts in the database that own tables: USER1 and USER2. Each user account has its own schema. Some examples for how the two users can access their own tables and tables owned by the other user follow:

USER1 accesses own TABLE1:	TABLE1
USER1 accesses own TEST:	TEST
USER1 accesses USER2's TABLE10:	USER2 . TABLE10
USER1 accesses USER2's TEST:	USER2 . TEST



**FIGURE 3.1**  
Schemas in a  
database.

In this example, both users have a table called TEST. Tables can have the same names in a database as long as they belong to different schemas. If you look at it this way, table names are always unique in a database because the schema owner is actually part of the table name. For instance, USER1 . TEST is a different table than USER2 . TEST. If you do not specify a schema with the table name when accessing tables in a database, the database server looks for a table that you own by default. That is, if USER1 tries to access TEST, the database server looks for a USER1-owned table named TEST before it looks for other objects owned by USER1, such as synonyms to tables in another schema. Hour 21, “Working with the System Catalog,” helps you fully understand how synonyms work. You must be careful to understand the distinction between objects in your own schema and those objects in another schema. If you do not provide a schema when performing operations that alter the table, such as a DROP command, the database will assume that you mean a table in your own schema. This could possibly lead to you unintentionally dropping the wrong object. So you must always pay careful attention as to which user you are currently logged into the database with.

Every database server has rules concerning how you can name objects and elements of objects, such as field names. You must check your particular implementation for the exact naming conventions or rules.

## A Table: The Primary Storage for Data

The table is the primary storage object for data in a relational database. In its simplest form, a table consists of row(s) and column(s), both of which hold the data. A table takes up physical space in a database and can be permanent or temporary.

### Columns

A field, also called a *column* in a relational database, is part of a table that is assigned a specific data type; a field should be named to correspond with the type of data that will be entered into that column. Columns can be specified as NULL or NOT NULL, meaning that if a column is NOT NULL, something must be entered. If a column is specified as NULL, nothing has to be entered.

Every database table must consist of at least one column. Columns are those elements within a table that hold specific types of data, such as a person's name or address. For example, a valid column in a customer table might be the customer's name. Figure 3.2 illustrates a column in a table.

**FIGURE 3.2**  
An Example of a  
Column

EMP_ID	LAST_NAME	FIRST_NAME	MIDDLE_NAME
213764555	GLASS	BRANDON	SCOTT
220984332	WALLACE	MARIAH	DAWN
311549802	STEPHENS	TINA	DAWN
313752439	GLASS	JACOB	DAWN
447346889	PLEW	LINDA	CAROL
443879912	SPURGEON	TIFFANY	DAWN

Generally, an object name must be one continuous string and can be limited to the number of characters used according to each implementation of SQL. It is typical to use underscores with names to provide separation between characters. For example, a column for the customer's name can be named CUSTOMER\_NAME instead of CUSTOMERNAME.

Additionally, data can be stored as either uppercase or lowercase for character-defined fields. The case that you use for data is simply a matter of preference, which should be based on how the data will be used. In many cases, data is stored in

uppercase for simplicity and consistency. However, if data is stored in different case types throughout the database (uppercase, lowercase, and mixed case), functions can be applied to convert the data to either uppercase or lowercase if needed. These functions will be covered in Hour 11, “Restructuring the Appearance of Data.”

Be sure to check your implementation for rules when naming objects and other database elements. Often database administrators will adopt a *naming convention* that explains how to name the objects within the database so you can easily discern how they are used.

**By the  
Way**

## Rows

A *row* is a record of data in a database table. For example, a row of data in a customer table might consist of a particular customer’s identification number, name, address, phone number, fax number, and so on. A row is comprised of fields that contain data from one record in a table. A table can contain as little as one row of data and up to as many as millions of rows of data or records. Figure 3.3 illustrates a row within a table.

EMP_ID	LAST_NAME	FIRST_NAME	MIDDLE_NAME
212764955	GLASS	BRANDON	SCOTT
220784122	WALLACE	MARSH	
311549802	STEPHENS	TINA	DAWN
312782439	GLASS	JACOB	
442346889	PLEW	LINDA	CAROL
442679912	SPURGEON	TIFFANY	

**FIGURE 3.3**  
Example of a  
Table Row.

## The CREATE TABLE Statement

The CREATE TABLE statement in SQL is used to create a table. Although the very act of creating a table is quite simple, much time and effort should be put into planning table structures before the actual execution of the CREATE TABLE statement. Carefully planning your table structure before implementation will save you from having to reconfigure things after they are in production.

Some elementary questions need to be answered when creating a table:

- ▶ What type of data will be entered into the table?
- ▶ What will be the table’s name?
- ▶ What column(s) will compose the primary key?

- ▶ What names shall be given to the columns (fields)?
- ▶ What data type will be assigned to each column?
- ▶ What will be the allocated length for each column?
- ▶ Which columns in a table can be left blank?

After these questions are answered, the actual CREATE TABLE statement is simple.

The basic syntax to create a table is as follows:

```
CREATE TABLE table_name
( field1 data_type [ not null ],
  field2 data_type [ not null ],
  field3 data_type [ not null ],
  field4 data_type [ not null ],
  field5 data_type [ not null ] );
```

A semicolon is the last character in the previous statement. Most SQL implementations have some character that terminates a statement or submits a statement to the database server. Oracle and MySQL use the semicolon. Transact-SQL has no such requirement. This book uses the semicolon.

**By the  
Way**

In this hour's examples, we use the popular data types CHAR (constant-length character), VARCHAR (variable-length character), NUMBER (numeric values, decimal and non-decimal), and DATE (date and time values).

Create a table called EMPLOYEE\_TBL in the following example:

```
CREATE TABLE EMPLOYEE_TBL
(EMP_ID      CHAR(9)      NOT NULL,
 EMP_NAME    VARCHAR (40)  NOT NULL,
 EMP_ST_ADDR VARCHAR (20)  NOT NULL,
 EMP_CITY    VARCHAR (15)  NOT NULL,
 EMP_ST      CHAR(2)      NOT NULL,
 EMP_ZIP     INTEGER(5)    NOT NULL,
 EMP_PHONE   INTEGER(10)   NULL,
 EMP_PAGER   INTEGER(10)   NULL);
```

Eight different columns make up this table. Notice the use of the underscore character to break the column names up into what appears to be separate words (EMPLOYEE ID is stored as EMP\_ID). This is a technique that is used to make a table or column name more readable. Each column has been assigned a specific data type and length, and by using the NULL/NOT NULL constraint, you have specified which columns require values for every row of data in the table. The EMP\_PHONE is defined as NULL, meaning that NULL values are allowed in this column because there might

be individuals without a telephone number. The information concerning each column is separated by a comma, with parentheses surrounding all columns (a left parenthesis before the first column and a right parenthesis following the information on the last column).

Each record, or row of data, in this table would consist of the following:

```
EMP_ID, EMP_NAME, EMP_ST_ADDR, EMP_CITY, EMP_ST, EMP_ZIP, EMP_PHONE, EMP_PAGER
```

In this table, each field is a column. The column EMP\_ID could consist of one employee's identification number or many employees' identification numbers, depending on the requirements of a database query or transactions. The column is a vertical entity in a table, whereas a row of data is a horizontal entity.

NULL is a default attribute for a column; therefore, it does not have to be entered in the CREATE TABLE statement. NOT NULL must always be specified.

**By the  
Way**

## Naming Conventions

When selecting names for objects, specifically tables and columns, the name should reflect the data that is to be stored. For example, the name for a table pertaining to employee information could be named EMPLOYEE\_TBL. Names for columns should follow the same logic. When storing an employee's phone number, an obvious name for that column would be PHONE\_NUMBER.

Check your particular implementation for name length limits and characters that are allowed; they could differ from implementation to implementation.

**By the  
Way**

## The ALTER TABLE Command

A table can be modified through the use of the ALTER TABLE command after that table's creation. You can add column(s), drop column(s), change column definitions, add and drop constraints, and, in some implementations, modify table STORAGE values. The standard syntax for the ALTER TABLE command follows:

```
alter table table_name [modify] [column column_name][datatype!null not null]
[restrict!cascade]
           [drop]   [constraint constraint_name]
           [add]   [column] column definition
```

## Modifying Elements of a Table

The *attributes* of a column refer to the rules and behavior of data in a column. You can modify the attributes of a column with the ALTER TABLE command. The word *attributes* here refers to the following:

- ▶ The data type of a column
- ▶ The length, precision, or scale of a column
- ▶ Whether the column can contain NULL values

The following example uses the ALTER TABLE command on EMPLOYEE\_TBL to modify the attributes of the column EMP\_ID:

```
ALTER TABLE EMPLOYEE_TBL MODIFY  
EMP_ID VARCHAR(10);
```

Table altered.

The column was already defined as data type VARCHAR (a varying-length character), but you increased the maximum length from 9 to 10.

## Adding Mandatory Columns to a Table

One of the basic rules for adding columns to an existing table is that the column you are adding cannot be defined as NOT NULL if data currently exists in the table. NOT NULL means that a column must contain some value for every row of data in the table. So, if you are adding a column defined as NOT NULL, you are contradicting the NOT NULL constraint right off the bat if the preexisting rows of data in the table do not have values for the new column.

There is, however, a way to add a mandatory column to a table:

1. Add the column and define it as NULL (the column does not have to contain a value).
2. Insert a value into the new column for every row of data in the table.
3. After ensuring that the column contains a value for every row of data in the table, you can alter the table to change the column's attribute to NOT NULL.

## Adding Auto-Incrementing Columns to a Table

Sometimes it is necessary to create a column that auto-increments itself in order to give a unique sequence number for a particular row. This could be done for many reasons, such as not having a natural key for the data or you would like to use a

unique sequence number to sort the data. Creating an auto-incrementing column is generally quite easy. In MySQL the implementation provides the SERIAL method to produce a truly unique value for the table. Following is an example:

```
CREATE TABLE TEST_INCREMENT(  
    ID SERIAL,  
    TEST_NAME VARCHAR(20));  
  
INSERT INTO TEST_INCREMENT(TEST_NAME)  
VALUES ('FRED'), ('JOE'), ('MIKE'), ('TED');  
  
SELECT * FROM TEST_INCREMENT;
```

ID	TEST_NAME
1	FRED
2	JOE
3	MIKE
4	TED

## Modifying Columns

There are many things to take into consideration when modifying existing columns of a table. Following are some common rules for modifying columns:

- ▶ The length of a column can be increased to the maximum length of the given data type.
- ▶ The length of a column can be decreased only if the largest value for that column in the table is less than or equal to the new length of the column.
- ▶ The number of digits for a number data type can always be increased.
- ▶ The number of digits for a number data type can be decreased only if the value with the most number of digits for that column is less than or equal to the new number of digits specified for the column.
- ▶ The number of decimal places for a number data type can either be increased or decreased.
- ▶ The data type of a column can normally be changed.

Some implementations may actually restrict you from using certain ALTER TABLE options. For example, you might not be allowed to drop columns from a table. To do this, you would have to drop the table itself, and then rebuild the table with the desired columns. You could run into problems by dropping a column in one table that is dependent on a column in another table, or a column that is referenced by a column in another table. Be sure to refer to your specific implementation documentation.

**Watch  
Out!**

Take heed when altering and dropping tables. If logical or typing mistakes are made when issuing these statements, important data can be lost.

## Creating a Table from an Existing Table

A copy of an existing table can be created using a combination of the CREATE TABLE statement and the SELECT statement. The new table has the same column definitions. Any or all columns can be selected. New columns that are created via functions or a combination of columns automatically assume the size necessary to hold the data. The basic syntax for creating a table from another table is as follows:

```
create table new_table_name as
select [ *|column1, column2 ]
from table_name
[ where ]
```

Notice some new keywords in the syntax, particularly the SELECT keyword. SELECT is a database query and is discussed in more detail in Chapter 7, “Introduction to Database Query.” However, it is important to know that you can create a table based on the results from a query.

First, we do a simple query to view the data in the PRODUCTS\_TBL table.

**Watch  
Out!**

You will create the tables that you see in these examples at the end of this hour in the “Exercises” section. In Hour 5, “Manipulating Data,” you will populate the tables you create in this hour with data.

```
select * from products_tbl;
```

PROD_ID	PROD_DESC	COST
11235	WITCHES COSTUME	29.99
222	PLASTIC PUMPKIN 18 INCH	7.75
13	FALSE PARAFFIN TEETH	1.1
90	LIGHTED LANTERNS	14.5
15	ASSORTED COSTUMES	10
9	CANDY CORN	1.35
6	PUMPKIN CANDY	1.45
87	PLASTIC SPIDERS	1.05
119	ASSORTED MASKS	4.95

**Watch  
Out!**

SELECT \* selects data from all fields in the given table. The \* represents a complete row of data, or record, in the table.



Next, create a table called PRODUCTS\_TMP based on the previous query:

```
create table products_tmp as  
select * from products_tbl;
```

Table created.

Now, if you run a query on the PRODUCTS\_TMP table, your results appear the same as if you had selected data from the original table.

```
select *  
from products_tmp;
```

PROD_ID	PROD_DESC	COST
11235	WITCHES COSTUME	29.99
222	PLASTIC PUMPKIN 18 INCH	7.75
13	FALSE PARAFFIN TEETH	1.1
90	LIGHTED LANTERNS	14.5
15	ASSORTED COSTUMES	10
9	CANDY CORN	1.35
6	PUMPKIN CANDY	1.45
87	PLASTIC SPIDERS	1.05
119	ASSORTED MASKS	4.95

When creating a table from an existing table, the new table takes on the same STORAGE attributes as the original table.

**By the  
Way**

## Dropping Tables

Dropping a table is actually one of the easiest things to do. When the RESTRICT option is used and the table is referenced by a view or constraint, the DROP statement returns an error. When the CASCADE option is used, the drop succeeds and all referencing views and constraints are dropped. The syntax to drop a table follows:

```
drop table table_name [ restrict|cascade ]
```

In the following example, you drop the table that you just created:

```
drop table products_tmp;
```

Table dropped.

Whenever dropping a table, be sure to specify the schema name or owner of the table before submitting your command. You could drop the incorrect table. If you have access to multiple user accounts, ensure that you are connected to the database through the correct user account before dropping tables.

**Watch  
Out!**

## Integrity Constraints

Integrity constraints are used to ensure accuracy and consistency of data in a relational database. Data integrity is handled in a relational database through the concept of referential integrity. Many types of integrity constraints play a role in referential integrity (RI).

### Primary Key Constraints

*Primary key* is the term used to identify one or more columns in a table that make a row of data unique. Although the primary key typically consists of one column in a table, more than one column can comprise the primary key. For example, either the employee's Social Security number or an assigned employee identification number is the logical primary key for an employee table. The objective is for every record to have a unique primary key or value for the employee's identification number.

Because there is probably no need to have more than one record for each employee in an employee table, the employee identification number makes a logical primary key. The primary key is assigned at table creation.

The following example identifies the EMP\_ID column as the PRIMARY KEY for the EMPLOYEES table:

```
CREATE TABLE EMPLOYEE_TBL
(EMP_ID      CHAR(9)      NOT NULL PRIMARY KEY,
EMP_NAME     VARCHAR (40)  NOT NULL,
EMP_ST_ADDR  VARCHAR (20)  NOT NULL,
EMP_CITY     VARCHAR (15)  NOT NULL,
EMP_ST       CHAR(2)      NOT NULL,
EMP_ZIP      INTEGER(5)   NOT NULL,
EMP_PHONE    INTEGER(10)  NULL,
EMP_PAGER    INTEGER(10)  NULL);
```

This method of defining a primary key is accomplished during table creation. The primary key in this case is an implied constraint. You can also specify a primary key explicitly as a constraint when setting up a table, as follows:

```
CREATE TABLE EMPLOYEE_TBL
(EMP_ID      CHAR(9)      NOT NULL,
EMP_NAME     VARCHAR (40)  NOT NULL,
EMP_ST_ADDR  VARCHAR (20)  NOT NULL,
EMP_CITY     VARCHAR (15)  NOT NULL,
EMP_ST       CHAR(2)      NOT NULL,
EMP_ZIP      INTEGER(5)   NOT NULL,
EMP_PHONE    INTEGER(10)  NULL,
EMP_PAGER    INTEGER(10)  NULL,
PRIMARY KEY (EMP_ID));
```

The primary key constraint in this example is defined after the column comma list in the CREATE TABLE statement.

A primary key that consists of more than one column can be defined by either of the following methods:

```
CREATE TABLE PRODUCTS
(PROD_ID      VARCHAR2(10)    NOT NULL,
 VENDOR_ID    VARCHAR2(10)    NOT NULL,
 PRODUCT      VARCHAR2(30)    NOT NULL,
 COST         NUMBER(8,2)     NOT NULL,
 PRIMARY KEY (PROD_ID, VENDOR_ID));

ALTER TABLE PRODUCTS
ADD CONSTRAINT PRODUCTS_PK PRIMARY KEY (PROD_ID, VENDOR_ID);
```

## Unique Constraints

A *unique column constraint* in a table is similar to a primary key in that the value in that column for every row of data in the table must have a unique value. Although a primary key constraint is placed on one column, you can place a unique constraint on another column even though it is not actually for use as the primary key.

Study the following example:

```
CREATE TABLE EMPLOYEE_TBL
(EMP_ID      CHAR(9)          NOT NULL      PRIMARY KEY,
 EMP_NAME    VARCHAR (40)     NOT NULL,
 EMP_ST_ADDR VARCHAR (20)     NOT NULL,
 EMP_CITY    VARCHAR (15)     NOT NULL,
 EMP_ST      CHAR(2)          NOT NULL,
 EMP_ZIP     INTEGER(5)       NOT NULL,
 EMP_PHONE   INTEGER(10)     NULL        UNIQUE,
 EMP_PAGER   INTEGER(10)     NULL);
```

The primary key in this example is EMP\_ID, meaning that the employee identification number is the column that is used to ensure that every record in the table is unique. The primary key is a column that is normally referenced in queries, particularly to join tables. The column EMP\_PHONE has been designated as a UNIQUE value, meaning that no two employees can have the same telephone number. There is not a lot of difference between the two, except that the primary key is used to provide an order to data in a table and, in the same respect, join related tables.

## Foreign Key Constraints

A *foreign key* is a column in a child table that references a primary key in the parent table. A *foreign key constraint* is the main mechanism used to enforce referential

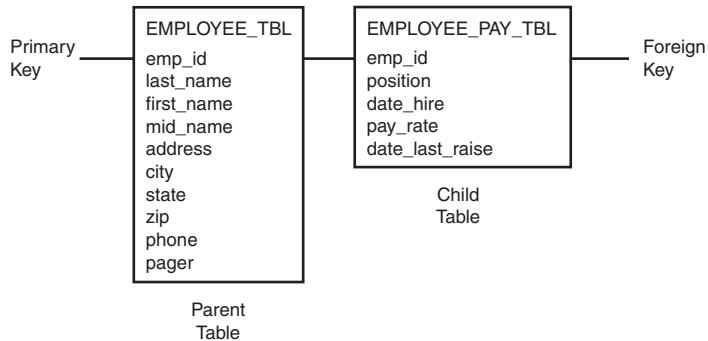
integrity between tables in a relational database. A column defined as a foreign key is used to reference a column defined as a primary key in another table.

Study the creation of the foreign key in the following example:

```
CREATE TABLE EMPLOYEE_PAY_TBL
(EMP_ID          CHAR(9)          NOT NULL,
 POSITION        VARCHAR2(15)     NOT NULL,
 DATE_HIRE      DATE              NULL,
 PAY_RATE       NUMBER(4,2)      NOT NULL,
 DATE_LAST_RAISE DATE            NULL,
 CONSTRAINT EMP_ID_FK FOREIGN KEY (EMP_ID) REFERENCES EMPLOYEE_TBL (EMP_ID));
```

The EMP\_ID column in this example has been designated as the foreign key for the EMPLOYEE\_PAY\_TBL table. This foreign key, as you can see, references the EMP\_ID column in the EMPLOYEE\_TBL table. This foreign key ensures that for every EMP\_ID in the EMPLOYEE\_PAY\_TBL, there is a corresponding EMP\_ID in the EMPLOYEE\_TBL. This is called a *parent/child relationship*. The parent table is the EMPLOYEE\_TBL table, and the child table is the EMPLOYEE\_PAY\_TBL table. Study Figure 3.4 for a better understanding of the parent table/child table relationship.

**FIGURE 3.4**  
The parent/child table relationship.



In this figure, the EMP\_ID column in the child table references the EMP\_ID column in the parent table. For a value to be inserted for EMP\_ID in the child table, a value for EMP\_ID in the parent table must first exist. Likewise, for a value to be removed for EMP\_ID in the parent table, all corresponding values for EMP\_ID must first be removed from the child table. This is how referential integrity works.

A foreign key can be added to a table using the ALTER TABLE command, as shown in the following example:

```
alter table employee_pay_tbl
add constraint id_fk foreign key (emp_id)
references employee_tbl (emp_id);
```

The options available with the ALTER TABLE command differ among different implementations of SQL, particularly when dealing with constraints. In addition, the actual use and definitions of constraints also vary, but the concept of referential integrity should be the same with all relational databases.

## NOT NULL Constraints

Previous examples use the keywords NULL and NOT NULL listed on the same line as each column and after the data type. NOT NULL is a constraint that you can place on a table's column. This constraint disallows the entrance of NULL values into a column; in other words, data is required in a NOT NULL column for each row of data in the table. NULL is generally the default for a column if NOT NULL is not specified, allowing NULL values in a column.

## Check Constraints

Check (CHK) constraints can be utilized to check the validity of data entered into particular table columns. Check constraints are used to provide back-end database edits, although edits are commonly found in the front-end application as well. General edits restrict values that can be entered into columns or objects, whether within the database itself or on a front-end application. The check constraint is a way of providing another protective layer for the data.

The following example illustrates the use of a check constraint:

```
CREATE TABLE EMPLOYEE_TBL
(EMP_ID          CHAR(9)          NOT NULL,
 EMP_NAME        VARCHAR2(40)     NOT NULL,
 EMP_ST_ADDR     VARCHAR2(20)     NOT NULL,
 EMP_CITY        VARCHAR2(15)     NOT NULL,
 EMP_ST          CHAR(2)          NOT NULL,
 EMP_ZIP         NUMBER(5)        NOT NULL,
 EMP_PHONE       NUMBER(10)       NULL,
 EMP_PAGER       NUMBER(10)       NULL),
PRIMARY KEY (EMP_ID),
CONSTRAINT CHK_EMP_ZIP CHECK ( EMP_ZIP = '46234' );
```

The check constraint in this table has been placed on the EMP\_ZIP column, ensuring that all employees entered into this table have a ZIP code of '46234'. Perhaps that is a little restricting. Nevertheless, you can see how it works.

If you wanted to use a check constraint to verify that the ZIP code is within a list of values, your constraint definition could look like the following:

```
CONSTRAINT CHK_EMP_ZIP CHECK ( EMP_ZIP in ('46234','46227','46745') );
```

If there is a minimum pay rate that can be designated for an employee, you could have a constraint that looks like the following:

```
CREATE TABLE EMPLOYEE_PAY_TBL
(EMP_ID          CHAR(9)          NOT NULL,
 POSITION         VARCHAR2(15)     NOT NULL,
 DATE_HIRE      DATE              NULL,
 PAY_RATE       NUMBER(4,2)       NOT NULL,
 DATE_LAST_RAISE DATE            NULL,
 CONSTRAINT EMP_ID_FK FOREIGN KEY (EMP_ID) REFERENCES EMPLOYEE_TBL (EMP_ID),
 CONSTRAINT CHK_PAY CHECK ( PAY_RATE > 12.50 ) );
```

In this example, any employee entered in this table must be paid more than \$12.50 an hour. You can use just about any condition in a check constraint, as you can with a SQL query. You learn more about these conditions in Hours 5 and 7.

## Dropping Constraints

Any constraint that you have defined can be dropped using the `ALTER TABLE` command with the `DROP CONSTRAINT` option. For example, to drop the primary key constraint in the `EMPLOYEES` table, you can use the following command:

```
ALTER TABLE EMPLOYEES DROP CONSTRAINT EMPLOYEES_PK;
```

Table altered.

Some implementations might provide shortcuts for dropping certain constraints. For example, to drop the primary key constraint for a table in MySQL, you can use the following command:

```
ALTER TABLE EMPLOYEES DROP PRIMARY KEY;
```

Table altered.

Some implementations allow you to disable constraints. Instead of permanently dropping a constraint from the database you might want to temporarily disable the constraint, and then enable it later.

## Summary

You have learned a little about database objects in general, but have specifically learned about the table. The table is the simplest form of data storage in a relational database. Tables contain groups of logical information, such as employee, customer, or product information. A table is composed of various columns, with each column having attributes; those attributes mainly consist of data types and constraints, such as `NOT NULL` values, primary keys, foreign keys, and unique values.

You learned the `CREATE TABLE` command and options, such as storage parameters, that might be available with this command. You have also learned how to modify the structure of existing tables using the `ALTER TABLE` command. Although the process of managing database tables might not be the most basic process in SQL, it is our philosophy that if you first learn the structure and nature of tables, you more easily grasp the concept of accessing the tables, whether through data manipulation operations or database queries. In later hours, you learn about the management of other objects in SQL, such as indexes on tables and views.

## Q&A

**Q.** *When I name a table that I am creating, is it necessary to use a suffix such as `_TBL`?*

**A.** Absolutely not. You do not have to use anything. For example, a table to hold employee information could be named similar to the following, or anything else that would refer to what type of data is to be stored in that particular table:

```
EMPLOYEE  
EMP_TBL  
EMPLOYEE_TBL  
EMPLOYEE_TABLE  
WORKER
```

**Q.** *Why is it so important to use the schema name when dropping a table?*

**A.** Here's a true story about a new DBA that dropped a table: A programmer had created a table under his schema with the same name as a production table. That particular programmer left the company. The programmer's database account was being deleted from the database, but the `DROP USER` statement returned an error due to the fact that outstanding objects were owned by the programmer. After some investigation, it was determined that the programmer's table was not needed, so a `DROP TABLE` statement was issued.

It worked like a charm—but the problem was that the DBA was logged in as the production schema when the `DROP TABLE` statement was issued. The DBA should have specified a schema name, or owner, for the table to be dropped. Yes, the wrong table in the wrong schema was dropped. It took approximately eight hours to restore the production database.

## Workshop

The following workshop is composed of a series of quiz questions and practical exercises. The quiz questions are designed to test your overall understanding of the current material. The practical exercises are intended to afford you the opportunity to apply the concepts discussed during the current hour, as well as build upon the knowledge acquired in previous hours of study. Please take time to complete the quiz questions and exercises before continuing. Refer to Appendix C, “Answers to Quizzes and Exercises,” for answers.

## Quiz

1. Will the following CREATE TABLE statement work? If not, what needs to be done to correct the problem(s)?

```

Create table EMPLOYEE_TABLE as:
( ssn          number(9)      not null,
  last_name    varchar2(20)   not null,
  first_name   varchar2(20)   not null,
  middle_name  varchar2(20)   not null,
  st address   varchar2(30)   not null,
  city         char(20)       not null,
  state        char(2)        not null,
  zip          number(4)      not null,
  date hired   date);

```

2. Can you drop a column from a table?
3. What statement would you issue in order to create a primary key constraint on the preceding EMPLOYEE\_TABLE?
4. What statement would you issue on the preceding EMPLOYEE\_TABLE to allow the MIDDLE\_NAME column to accept NULL values?
5. What statement would you use to restrict the people added into the preceding EMPLOYEE\_TABLE to only reside in the state of New York ('NY')?
6. What statement would you use to add an auto-incrementing column called EMPID to the preceding EMPLOYEE\_TABLE?



## Exercises

1. Bring up a command prompt and use the following syntax to log onto your local MySQL instance, replacing *username* with your username and *password* with your password. Ensure that you do not leave a space between `-p` and your password.

```
mysql -h localhost -u username -ppassword
```

2. At the `mysql>` command prompt, enter the following command to tell MySQL that you want to use the database you created previously:

```
use learnsql;
```

3. Now, go to Appendix D, “CREATE TABLE Statements for Book Examples,” to get the DDL for the tables used in this book. At the `mysql>` prompt, enter each CREATE TABLE statement. Be sure to include a semicolon at the end of each CREATE TABLE statement. The tables that you create will be used throughout the book.

4. At the `mysql>` prompt, enter the following command to get a list of your tables:

```
show tables;
```

5. At the `mysql>` prompt, use the DESCRIBE command (`desc` for short) to list the columns and their attributes for each one of the tables you created. For example:

```
describe employee_tbl;  
describe employee_pay_tbl;
```

6. If you have any errors or typos, simply re-create the appropriate table(s). If the table was successfully created, but has typos (perhaps you did not properly define a column or forgot a column), drop the table, and issue the CREATE TABLE command again. The syntax of the DROP TABLE command is as follows:

```
drop table orders_tbl;
```

# Index

## Symbols

- + (addition operator), 135, 210
- / (division operator), 136
- || (double pipe signs), 353
- = (equal operator), 118
- > (greater than operator), 119-120
- < (less than operator), 119-120
- \* (multiplication operator), 135-136
- != (non-equality operator), 119
- ;(semicolons), 46
- " (single quotation marks), 353
- (subtraction operator), 135

## A

- abandoned privileges, 304
- ABS (absolute value) function, 178
- accessing
  - remote databases, 361
    - JDBC, 363
    - ODBC, 362

- vendor connectivity tools, 363
- web interface, 363-364
- user access, controlling, 302, 361
  - columns, 304
  - GRANT statement, 302-303
  - groups of privileges, 305
  - PUBLIC database account, 304
  - REVOKE statement, 303-304
- adding
  - auto-incrementing columns to tables, 48
  - characters to strings, 176-177
  - columns to tables, 48
  - data to tables, 74-75
    - from another table, 76-78
    - NULL values, 78-79
    - into specified columns, 75-76
  - rows into views, 321
  - time to dates, 190-191
- addition operator (+), 135
- ADD\_MONTHS function, 190
- ADMIN OPTION (GRANT statement), 303

**aggregate functions****aggregate functions**

AVG, 146-147  
 COUNT, 142-144  
 definition, 141-142  
 GROUP BY clause, 153-156  
 MAX, 147  
 MIN, 147-148  
 SUM, 144-146

**aliases**

columns, 112-113  
 tables, 208

**ALL operator, 126****ALL option (SELECT statement), 104****ALTER TABLE statement, 47-48, 381****American National Standards Institute. See ANSI****AND operator, 127-128****ANSI (American National Standards Institute), 8**

character functions, 165  
   concatenation, 166-168  
   INSTR, 172  
   LOWER, 170  
   LTRIM, 173  
   REPLACE, 169  
   RTRIM, 173-174  
   SUBSTR, 170-171  
   substrings, 166  
   TRANSLATE, 166-169  
   UPPER, 169  
 object privileges, 300  
 SELECT statement syntax, 370  
 trigger creation syntax, 349

**ANSI SQL, 8, 371****ANY operator, 126****arithmetic operators, 134**

addition, 135  
 combining, 136-137

division, 136  
 multiplication, 135-136  
 subtraction, 135

**ascending order, 106****ASCII characters, returning, 178****ASCII chart website, 178****ASCII function, 178****authIDs (Authorization Identifiers), 283****authority levels, 305****AUTHORIZATION keyword (CREATE SCHEMA statement), 290****auto-incrementing columns, 48****automated population, 74****AVG function, 146-147****avoiding**

indexes, 259-260  
 large sort operations, 275

**B****back-end applications, 360-361****base tables, join considerations, 214-215****BETWEEN operator, 122, 222****BLOB data type, 30****book website, 9****BOOLEAN data types, 34****C****call-level interface (CLI), 352****Cartesian product, 215-217****CASCADE option (REVOKE statement), 303**

- case sensitivity (queries), 108**
- CEIL function, 178**
- ceiling values function, 178**
- Center for Internet Security website, 298**
- CHAR data type, 29**
- character functions, 165**
  - ASCII, 178
  - COALESCE, 176
  - combining, 181-182
  - concatenation, 166-168
  - DECODE, 174-175
  - IFNULL, 175-176
  - INSTR, 172
  - LENGTH, 175
  - LOWER, 170
  - LPAD, 176-177
  - LTRIM, 173
  - REPLACE, 169
  - RPAD, 177
  - RTRIM, 173-174
  - SUBSTR, 170-171
  - substrings, 166
  - TRANSLATE, 166-169
  - UPPER, 169
- character string conversions**
  - dates, 196-197
  - to numbers, 179-180
- characters**
  - adding to strings, 176-177
  - ASCII, returning, 178
  - constant, 29
  - lowercase, 170
  - positions, 172
  - replacing, 169
  - trimming, 173-174
  - uppercase, 169
- CHK (check) constraints, 55-56**
- clauses**
  - FROM, 385
    - SELECT statement, 104
    - table arrangement, 269
  - GROUP BY, 152, 385
    - aggregate functions, 153-156
    - compared to ORDER BY clause, 156-159
    - compound queries, 244-245
    - CREATE VIEW statement, 323
    - functions, 152
    - ordering column names with numbers, 156
    - selected data, 152
  - HAVING, 159-160, 275, 385
  - ORDER BY, 385
    - compared to GROUP BY clause, 156-159
    - compound queries, 242-244
    - SELECT statement, 106-108
    - views, 323
  - SELECT, 102, 384
  - WHERE, 385
    - DELETE statement, 81
    - restrictive condition, 270-271
    - SELECT statement, 105-106
- CLI (call-level interface), 352**
- client/server systems, 12**
- closing cursors, 345-346**
- COALESCE function, 176**
- Codd, Dr. E.F., 8**
- columns, 21, 44-45**
  - adding, 48
  - aliases, 112-113
  - attributes, editing, 48

**columns**

- auto-incrementing, adding, 48
- averaging values, 146-147
- cardinality, 260
- check constraints, 55-56
- counting values, 142-144
- data, adding, 75-76
- dropping constraints, 56
- editing, 49
- foreign keys, 53-54
- index considerations, 258
- maximum values, 147
- minimum values, 147-148
- NOT NULL constraints, 55
- NULL values, 78-79
- ordering with numbers, 156
- primary keys, 52-53
- qualifying, 205
- totaling values, 144-146
- unique constraints, 53
- updating, 79-80
- user access control, 304

**combining**

- arithmetic operators, 136-137
- character functions, 181-182
- comparison operators, 120-121

**commands. See statements****COMMIT statement, 89-90, 381****comparison operators, 118**

- combining, 120-121
- equal, 118
- less than, greater than, 119-120
- non-equality, 119

**composite indexes, 257****compound queries, 235**

- clauses
  - GROUP BY, 244-245
  - ORDER BY, 242-244
- data retrieval, 246
- operators
  - EXCEPT, 241-242
  - INTERSECT, 240-241
  - UNION, 237-240

**concatenation, 166-168****conditions, queries, 105-106****conjunctive operators, 127**

- AND, 127-128
- OR, 128-130

**CONNECT statement, 14****CONNECT group, 305****connecting sessions, 14****constant characters, 29****constraints (integrity), 52**

- check, 55-56
- dropping, 56
- foreign keys, 53-54
- NOT NULL, 55
- primary keys, 52-53
- unique, 53

**controlling**

- data, 16
- transactions, 88-89
  - COMMIT statement, 89-90
  - performance, 95
  - RELEASE SAVEPOINT statement, 94
  - ROLLBACK statement, 90-92
  - ROLLBACK TO SAVEPOINT statement, 92-94
  - SAVEPOINT statement, 92
  - SET TRANSACTION statement, 94
- statements, 17

- user access, 302
  - columns, 304
  - GRANT statement, 302-303
  - groups of privileges, 305
  - PUBLIC database account, 304
  - REVOKE statement, 303-304
- conversion functions, 179**
  - character strings to numbers, 179-180
  - numeric strings to characters, 180-181
- converting dates, 192**
  - character strings, 196-197
  - date pictures, 193-195
- correlated subqueries, 229-230**
- COUNT function, 111, 142-144**
- counting table records, 111**
- CREATE DOMAIN statement, 381**
- CREATE INDEX statement, 255, 381**
- CREATE ROLE statement, 306, 382**
- CREATE SCHEMA statement, 289-290**
- CREATE TABLE AS statement, 382**
- CREATE TABLE statement, 45-47, 50-51, 382**
  - CUSTOMER TBL statement, 436
  - EMPLOYEE PAY TBL statement, 435
  - EMPLOYEE TBL statement, 435
  - ORDERS TBL statement, 436
  - PRODUCTS TBL statement, 436
- CREATE TRIGGER statement, 349-350**
- CREATE TYPE statement, 382**
- CREATE VIEW statement, 316, 382**
  - GROUP BY clause, 323
  - views from multiple tables, 318-319
  - views from other views, 319-320
  - views from single tables, 316-318
  - WITH CHECK OPTION, 320-321
- creating**
  - indexes, 255
  - roles, 306
  - schemas, 289-290
  - SQL with SQL, 352-353
  - synonyms, 324-325
  - system catalog, 331
  - tables, 45-47
    - existing tables, 50-51
    - from views, 322
  - triggers, 349-350
  - users, 286
    - MySQL, 289
    - Oracle, 287-288
    - SQL Server, 288-289
    - Sybase, 288-289
  - views
    - from single tables, 316-318
    - from multiple tables, 318-319
    - from other views, 319-320
    - WITH CHECK OPTION, 320-321
- cross joins, 215-217**
- current date/time function, 188**
- cursors**
  - closing, 345-346
  - current values, 344
  - declaring, 344
  - definition, 343
  - fetching data from, 345
  - opening, 345
  - overview, 344

**data****D****data**

- administration, 17
- controlling, 16
- definition, 27
- fetching from cursors, 345
- for indexes, 272
- grouping, 151
  - GROUP BY clause, 152-156
  - GROUP BY clause versus ORDER BY clause, 156-159
  - HAVING clause, 159-160
- manipulating, 16, 73
- populating tables, 74
- redundancy, 63
- retrieving from compound queries, 246
- selecting
  - statements, 16
  - multiple tables, 203
- simplifying with views, 314
- summarized data maintenance, 315-316
- system catalog, 331-332
- tables
  - deleting, 81
  - examples in book, 18-20
  - inserting, 74-75
  - inserting from another table, 76-78
  - inserting into specified columns, 75-76
  - inserting NULL values, 78-79
  - selecting from another table, 112
  - updating, 79-80
- views, updating, 321

**Data Control Language (DCL), 16****Data Definition Language (DDL), 15****data dictionaries. See system catalog****Data Manipulation Language. See DML****Data Query Language (DQL), 16****data types**

- basic, 28
- BLOB, 30
- BOOLEAN, 34
- CHAR, 29
- date and time, 32-33, 186-187
- decimal, 31-32
- definition, 27
- domains, 35
- DOUBLE PRECISION, 32
- fixed-length strings, 29
- floating-point decimals, 32
- integers, 32
- large objects, 30
- lengths, 37
- literal strings, 33-34
- NULL, 34
- numeric, 30-31
- REAL, 32
- TEXT, 30
- user-defined, 35
- VARCHAR, 29
- varying-length strings, 29

**database administrators (DBAs), 285****database management system (DBMS), 7****databases**

- client/server systems, 12
- definition, 10
- denormalizing, 69
- design information, 332
- full table scans, 254
- Internet access tools, 365

- logical, 62-63
- MySQL examples/exercises, 22
- normalizing
  - benefits, 67-68
  - disadvantages, 68
  - names, 67
  - normal forms, 61, 64-66
  - overview, 61-62
- objects
  - definition, 41
  - schemas, 42-43
- parsing, 275
- queries. *See also* subqueries
  - case sensitivity, 108
  - column aliases, 112-113
  - compound. *See* compound queries
  - conditions, 105-106
  - counting table records, 111
  - definition, 16, 101
  - examples, 109-110
  - grouping results. *See* groups, data
  - ordering output, 106-108
  - searching, 174-175
  - SELECT statement, 101-104
  - SELECT statement with case sensitivity, 108
  - SELECT statement with FROM clause, 104
  - SELECT statement with ORDER BY clause, 106-108
  - SELECT statement with WHERE clause, 105-106
  - selecting data from another table, 112
  - single, 235
- raw, 62
- relational, 11
- remote, 361-364
- security, 297-298
  - privileges. *See* privileges
  - user access control, 302-305
- structures statements, 15
- transactions
  - statements, 17
  - controlling, 88-90
  - definition, 87
  - initiating, 94
  - overview, 87
  - performance, 95
  - savepoints, 92-94
  - saving changes, 89-90
  - undoing, 90-92
- tuning, 266
- users
  - authIDs, 283
  - creating, 286-287
  - creating in MySQL, 289
  - creating in Oracle, 287-288
  - creating in SQL Server, 288-289
  - creating in Sybase, 288-289
  - deleting, 293
  - editing, 291
  - GUI tools, 293
  - managing, 285
  - roles/privileges, 285
  - schemas, 286-290
  - sessions, 292
  - types, 284
- vendors, 13-14
- web-based systems, 12-13
- date and time data types, 32-33**



**DATEADD function****DATEADD function, 190****DATEDIFF function, 192****DATENAME function, 192****DATEPART function, 192****dates**

conversions, 192

character strings, 196-197

date pictures, 193-195

data types

implementation-specific, 187

standard, 186

date functions, 187

adding time, 190-191

comparing dates/times, 191

current, 188

miscellaneous, 192

time zones, 189

DATETIME elements, 186

parts, 194-195

pictures, 193-195

storing, 186

system, 188

**DATETIME data types, 32****DATETIME element, 186****DAYNAME function, 192****DAYOFMONTH function, 192****DAYOFWEEK function, 192****DAYOFYEAR function, 192****DBA group, 305****DBAs (database administrators), 285****dBASE, 333****DBMS (database management system), 7****DCL (Data Control Language), 16****DDL (Data Definition Language), 15****DECIMAL data type, 31****decimals, 31-32****DECODE function, 174-175****DELETE statement, 383**

subqueries, 226

table data, 81

WHERE clause, 81

**deleting**

rows into views, 321

savepoints, 94

schemas, 290

table data, 81

users, 293

**denormalization, 69****descending order, 106****differences in vendor implementations, 369-371****direct SQL, 353****DISCONNECT statement, 14****disconnecting sessions, 14****DISTINCT statement, 104, 142****division operator (/), 136****DML (Data Manipulation Language), 16**

DELETE statement

deleting table data, 81

subqueries, 226

INSERT statement

adding data from another table, 76-78

adding data to specific columns, 75-76

adding data to tables, 74-75

subqueries, 224-225

NULL values, 78-79

overview, 73

UPDATE statement

multiple columns, 80

single columns, 79-80

- subqueries, 225-226
- tables, 79
- domain data types, 35
- double pipe signs (||), 353
- DOUBLE PRECISION data type, 32
- DQL (Data Query Language), 16
- DROP statement, 51
  - indexes, 260-261
  - users, 293
- DROP INDEX statement, 383
- DROP ROLE statement, 306
- DROP SCHEMA statement, 290
- DROP TABLE statement, 383
- DROP TRIGGER statement, 351
- DROP VIEW statement, 323, 383
- dropping
  - constraints, 56
  - indexes, 260-261
  - roles, 306
  - synonyms, 325
  - tables, 51, 57
  - triggers, 351
  - views, 323
- dynamic SQL, 351-352

## E

- editing
  - columns, 49
  - tables, 47-49
  - users, 291
- embedded SQL, 353
- embedding subqueries, 227-228

- enhancements, 371
- enterprise, 359-361
- equal operator (=), 118
- equijoins, 204-206
- example extensions, 372-373
  - MySQL, 374-375
  - PL/SQL, 373-374
  - Transact-SQL, 373
- EXCEPT operator (compound queries), 241-242
- EXISTS operator, 125
- EXIT statement, 14
- exiting sessions, 14
- EXP (exponential values) function, 178
- extensions, 371-372
  - MySQL, 374-375
  - PL/SQL, 373-374
  - Transact-SQL, 373

## F

- FETCH statement, 345
- fetching data from cursors, 345
- fields (tables), 20
- firewalls, 364
- first normal forms, 64
- fixed-length strings, 29
- FLOAT data type, 32
- floating-point decimals, 32
- FLOOR function, 178
- floor values function, 178
- FOR EACH ROW syntax (triggers), 351
- foreign keys, 53-54
- forgotten passwords, 308

**formatting statements****formatting statements, 266**

- FROM clause table arrangement, 269
- join order, 269-270
- readability, 267-269
- WHERE clause condition, 270-271

**FROM clause, 385**

- SELECT statement, 104
- table arrangement, 269

**front-end applications, 360-361****front-end tools, 63****full table scans, 254, 272****functions**

- ADD\_MONTHS, 190
- aggregate
  - AVG, 146-147
  - COUNT, 142-144
  - definition, 141-142
  - GROUP BY clause, 153-156
  - MAX, 147
  - MIN, 147-148
  - SUM, 144-146
- character, 165
  - ASCII, 178
  - COALESCE, 176
  - combining, 181-182
  - concatenation, 166-168
  - DECODE, 174-175
  - IFNULL, 175-176
  - INSTR, 172
  - LENGTH, 175
  - LOWER, 170
  - LPAD, 176-177
  - LTRIM, 173
  - REPLACE, 169

RPAD, 177

RTRIM, 173-174

SUBSTR, 170-171

substrings, 166

TRANSLATE, 166-169

UPPER, 169

## conversion

- character strings to numbers, 179-180
- numeric strings to characters, 180-181

COUNT, 111

date, 187

- adding time, 190-191
- comparing dates/times, 191
- current, 188
- miscellaneous, 192
- time zones, 189

DATEADD, 190

DATEDIFF, 192

DATENAME, 192

DATEPART, 192

DAYNAME, 192

DAYOFMONTH, 192

DAYOFWEEK, 192

DAYOFYEAR, 192

definition, 141, 347

GETDATE(), 188, 192

GROUP BY clause, 152

mathematical, 178

MONTHS\_BETWEEN, 192

NEXT\_DAY, 192

NOW, 188

**G****GETDATE() function, 188, 192****GRANT statement, 383**

ADMIN OPTION, 303

GRANT OPTION, 303

privileges, 301

user access control, 302-303

**granting privileges, 301****greater than operator ( > ), 119-120****GROUP BY clause, 385**

aggregate functions, 153-156

compared to ORDER BY clause, 156-159

compound queries, 244-245

CREATE VIEW statement, 323

functions, 152

ordering column names with numbers, 156

selected data, 152

**groups**

data, 151

GROUP BY clause, 152-156

GROUP BY clause versus ORDER BY clause,  
156-159

HAVING clause, 159-160

privileges, 305

**GUI tools, 293****H - I****HAVING clause, 159-160, 275, 385****IFNULL function, 175-176****implementation-specific data types, 187****implementations**

ANSI SQL compliance, 371

cursors, 344

differences, 369-371

extensions, 371

SQL, 10

system catalog, 333-334

**implicit indexes, 257****IN operator, 123****indexes**

avoiding, 259-260

column considerations, 258

creating, 255

data for, 272

definition, 253

disabling during batch loads, 275-276

dropping, 260-261

function, 254-255

overview, 253-254

performance, 260, 275-276

types, 255

composite, 257

implicit, 257

single-column, 256

unique, 256-258

**Informix, 371****initiating transactions, 94****INSERT object privilege, 300****INSERT statement, 383**

adding data to tables, 74

from another table, 76-78

specified columns, 75-76

CUSTOMER TBL statement, 438

EMPLOYEE PAY TBL statement, 438

**INSERT statement**

EMPLOYEE TBL statement, 437  
 ORDERS TBL statement, 439  
 PRODUCTS TBL statement, 440  
 subqueries, 224-225

**INSERT(column\_name) object privilege, 300**

**INSERT...SELECT statement, 383**

**installing MySQL**

Linux, 388-389  
 Windows, 387-388

**INSTR function, 172**

**integers, 32****integrity constraints, 52**

check, 55-56  
 dropping, 56  
 foreign keys, 53-54  
 NOT NULL, 55  
 primary keys, 52-53  
 unique, 53

**interactive SQL statements, 375-376**

**International Standards Organization (ISO), 8**

**Internet**

data availability for employees/customers, 365  
 database access tools, 365  
 security, 366  
 worldwide availability, 364

**INTERSECT operator (compound queries), 240-241**

**intranets, 365-366**

**INX suffix, 18**

**IS NOT NULL operator, 133**

**IS NULL operator, 121-122**

**ISO (International Standards Organization), 8**

**J**

**JDBC (Java Database Connectivity), 363**

**joins**

base tables, 214-215  
 Cartesian product, 215-217  
 component locations, 204  
 equijoins, 204-206  
 multiple keys, 213-214  
 natural, 206-207  
 non-equijoins, 208-209  
 ordering, 269-270  
 outer, 210-211  
 self, 212-213  
 table aliases, 208  
 types, 204

**K-L****keys**

foreign, 53-54  
 joining, 213-214  
 primary, 21, 52-53

**large object data types, 30**

**LENGTH function, 175**

**lengths**

data types, 37  
 strings, 175

**less than operator (<), 119-120**

**LIKE operator, 123-124, 273**

**Linux, MySQL installation, 388-389**

**literal strings, 33-34**

**logical databases, 62-63**

**logical operators, 121**

- ALL, 126
- ANY, 126
- BETWEEN, 122
- EXISTS, 125
- IN, 123
- IS NULL, 121-122
- LIKE, 123-124
- SOME, 126

**LOWER function, 170****lowercase strings, 170****LPAD function, 176-177****LTRIM function, 173****M****managing users, 285**

- creating users, 286-287
  - MySQL, 289
  - Oracle, 287-288
  - SQL Server, 288-289
  - Sybase, 288-289
- deleting, 293
- editing, 291
- GUI tools, 293
- schemas, 289-290
- sessions, 292

**manipulating data, 16, 73****manual population of data, 74****mathematical functions, 178****MAX function, 147****Microsoft**

- Access, 333
- SQL Server, users, 288-289

**MIN function, 147-148****MONTHS\_BETWEEN function, 192****multiplication operator (\*), 135-136****MySQL, 374-375**

- cursor declaration, 344
- examples/exercises, 22
- installing
  - Linux, 388-389
  - Windows, 387-388
- stored procedure syntax, 347-348
- system catalog implementations, 334
- system privileges, 300
- trigger creation syntax, 350
- users, creating, 289
- website, 375

**N****names**

- normalization, 67
- saving points, 92
- synonyms, 326
- tables, 18, 47

**natural joins, 206-207****negative operators, 130**

- IS NOT NULL, 133
- NOT BETWEEN, 131-132
- not equal, 131
- NOT EXISTS, 134
- NOT IN, 132
- NOT LIKE, 133

**nesting**

- queries. See subqueries
- stored procedures, 346

**Net8****Net8, 363****NEXT\_DAY function, 192****non-equality operator (!=), 119****non-equijoins, 208-209****normal forms, 61, 64**

first, 64

second, 65

third, 66

**normalization**

benefits, 67-68

disadvantages, 68

names, 67

normal forms, 61

first, 64

second, 65

third, 66

overview, 61-62

**NOT BETWEEN operator, 131-132****NOT EXISTS operator, 134****NOT IN operator, 132****NOT LIKE operator, 133****NOT NULL constraints, 55****NOW function, 188****NULL data types, 34****NULL value checker, 175-176****NULL values**

adding to columns, 78-79

checking, 175-176

replacing, 176

tables, 22

**NUMERIC data type, 30-31****numeric strings, converting to characters, 180-181****O****object privileges, 300-301****ODBC (Open Database Connectivity), 362****Open Client/C Developers Kit, 363****opening cursors, 345****operators**

arithmetic, 134

addition, 135

combining, 136-137

division, 136

multiplication, 135-136

subtraction, 135

**BETWEEN, 222**

comparison

combining, 120-121

equal, 118

less than, greater than, 119-120

non-equality, 119

conjunctive

AND, 127-128

OR, 128-130

definition, 105, 117

**EXCEPT, 241-242****INTERSECT, 240-241****LIKE, 273**

logical

ALL, 126

ANY, 126

BETWEEN, 122

EXISTS, 125

IN, 123

IS NULL, 121-122

LIKE, 123-124

SOME, 126

- negative, 130
    - IS NOT NULL, 133
    - NOT BETWEEN, 131-132
    - not equal, 131
    - NOT EXISTS, 134
    - NOT IN, 132
    - NOT LIKE, 133
  - OR, 274-275
  - OVERLAPS, 191
  - UNION, 235-239
  - UNION ALL, 239-240
  - options**
    - ADMIN OPTION, 303
    - ALL, 104
    - CASCADE, 303
    - DISTINCT, 104
    - GRANT OPTION, 303
    - RESTRICT, 303
    - WITH CHECK, 320-321
  - OR operator, 128-130, 274-275**
  - Oracle**
    - cursor declaration, 344
    - Net8, 363
    - parameters, 376
    - PL/SQL, 373-374
    - roles, 305
    - SELECT statement syntax, 370
    - stored procedure syntax, 347-348
    - system catalog implementations, 334
    - system privileges, 299
    - trigger creation syntax, 350
    - users, creating, 287-288
  - ORDER BY clause, 385**
    - compared to GROUP BY clause, 156-159
    - compound queries, 242-244
    - SELECT statement, 106-108
    - views, 323
  - outer joins, 210-211**
  - OVERLAPS operator, 191**
  - owners (schemas), 42**
- ## P
- parameters, 375**
  - parent/child table relationships, 54**
  - parsing, 275**
  - parts of dates, 194-195**
  - passwords**
    - forgotten, 308
    - system catalog, 338
  - performance**
    - definition, 265-266
    - formatting, 266
    - FROM clause table arrangement, 269
    - full table scans, 272
    - HAVING clause, 275
    - indexes, 260, 275-276
    - join order, 269-270
    - large sort operations, 275
    - LIKE operator, 273
    - OR operator, 274-275
    - readability, 267-269
    - stored procedures, 275
    - statistics stored in system catalog, 332
    - tools, 276
    - transactional control, 95
    - WHERE clause condition, 270-271
    - wildcard placement, 273



**PL/SQL****PL/SQL, 373-374**

plus (+) symbol, 210

**populating tables with data, 74-75**

from another table, 76-78

NULL values, 78-79

into specified columns, 75-76

**positioning characters, 172****POWER function, 178**

precision, 31

primary keys, 21, 52-53

**PRIVATE synonyms, 324****privileges, 298**

abandoned, 304

controlling with roles, 305-307

granting/revoking, 301

groups, 305

object, 300-301

system, 299-300

**pseudocolumns, 188****PUBLIC database account, 304****PUBLIC synonyms, 324**

UNION ALL operator, 239-240

UNION operator, 237-239

conditions, 105-106

counting table records, 111

definition, 16, 101

examples, 109-110

grouping results, 151

GROUP BY clause, 152-156

GROUP BY clause versus ORDER BY clause,  
156-159

HAVING clause, 159-160

ordering output, 106-108

searching, 174-175

SELECT statement, 101

case sensitivity, 108

FROM clauses, 104

ORDER BY clauses, 106-108

selecting data, 102-104

WHERE clauses, 105-106

selecting data from another table, 112

single, 235

system catalog, 335-336

**Q****qualifying columns, 205****queries. See also subqueries, 221**

case sensitivity, 108

column aliases, 112-113

compound, 235

data retrieval, 246

EXCEPT operator, 241-242

GROUP BY clause, 244-245

INTERSECT operator, 240-241

ORDER BY clause, 242-244

**R****raw databases, 62****RDBMS (relational database management system), 7****readability of statements, 267-269****REAL data type, 32**

records (tables), 21, 111

redundancy (data), 63

REFERENCES object privilege, 301

**REFERENCES(column\_name) object privilege, 301**

referential integrity, 68

relational database management system  
(RDBMS), 7

relational databases, 11

**RELEASE SAVEPOINT statement, 94**

remote databases, accessing, 361

JDBC, 363

ODBC, 362

vendor connectivity tools, 363

web interface, 363-364

**REPLACE function, 169**

replacing

characters, 169

NULL values, 176

**RESOURCE group, 305**

**RESTRICT keyword**

DROP SCHEMA statement, 290

REVOKE statement, 303

**REVOKE statement, 384**

privileges, 301

user access control, 303-304

users, 293

revoking privileges, 301

roles

creating, 306

dropping, 306

Oracle, 305

setting, 307

**ROLLBACK statement, 90-92, 384**

**ROLLBACK TO SAVEPOINT statement, 92-94**

rolling back savepoints, 92-94

**ROUND function, 178**

**rows, 21, 45**

averaging values, 146-147

counting, 142-144

maximum values, 147

minimum values, 147-148

totaling values, 144-146

views, 321

**RPAD function, 177**

**RTRIM function, 173-174**

## S

**SAVEPOINT statement, 92, 384**

savepoints

deleting, 94

names, 92

rolling back, 92-94

schemas

creating, 289-290

definition, 42

deleting, 290

overview, 42-43

owners, 42

users, compared, 286

**searching queries, 174-175**

**second normal forms, 65**

security

databases, 297-298

firewalls, 364

information stored in system catalog, 332

Internet, 366

privileges, 298

abandoned, 304

controlling with roles, 305-307

**security**

- granting/revoking, 301
- groups, 305
- object, 300-301
- system, 299-300
- roles, 305
  - creating, 306
  - dropping, 306
  - setting, 307
- user access
  - columns, 304
  - GRANT statement, 302-303
  - groups of privileges, 305
  - PUBLIC database account, 304
  - REVOKE statement, 303-304
- views, 315
- security officers, 285**
- SELECT statement, 384**
  - clauses, 102
  - column aliases, 112-113
  - COUNT function, 111
  - EXCEPT operator, 241-242
  - GROUP BY clause, 244-245
    - aggregate functions, 153-156
    - compared to ORDER BY clause, 156-159
    - functions, 152
    - ordering column names with numbers, 156
    - selected data, 152
  - HAVING clause, 159-160
  - implementation differences, 370
  - INTERSECT operator, 240-241
  - ORDER BY clause, 242-244
  - queries, 101
    - ALL option, 104
    - case sensitivity, 108
    - DISTINCT option, 104
    - FROM clause, 104
    - ORDER BY clause, 106-108
    - selecting data, 102-104
    - WHERE clause, 105-106
  - selecting data from another table, 112
  - single queries, 235
  - subqueries, 223-224
  - UNION ALL operator, 239-240
  - UNION operator, 237-239
- SELECT object privilege, 300**
- selecting data**
  - from another table, 112
  - statements, 16
  - multiple tables, 203
- self joins, 212-213**
- semicolons (;), 46**
- sessions**
  - connecting, 14
  - definition, 14
  - disconnecting, 14
  - exiting, 14
  - users, 292
- SET ROLE statement, 307**
- SET TRANSACTION statement, 94**
- SIGN function, 178**
- sign values function (SIGN), 178**
- single queries, 235**
- single quotation marks ("), 353**
- single-column indexes, 256**
- SOME operator, 126**
- sort operations, 275**
- SQL (Structured Query Language), 8**
  - definition, 8
  - generation with SQL, 352-353
  - implementation, 10

- on the Internet
  - data availability for employees/customers, 365
  - database access tools, 365
  - worldwide availability, 364
- optimizer, 267
- SQL Server**
  - cursor declaration, 344
  - stored procedure syntax, 347-348
  - system catalog implementations, 333
  - Transact-SQL, 373
  - trigger creation syntax, 350
  - users, creating, 288-289
- SQL-2003, 9-10**
- SQLBase**
  - authority levels, 305
  - SELECT statement syntax, 370
- SQRT (square root) function, 178**
- standard data types, 186**
- standards**
  - ANSI SQL, 8
  - SQL-2003, 9-10
  - table-naming, 18
- statements**
  - ALTER TABLE, 47-48, 381
  - COMMIT, 89-90, 381
  - CONNECT, 14
  - CREATE DOMAIN, 381
  - CREATE INDEX, 255, 381
  - CREATE ROLE, 306, 382
  - CREATE SCHEMA, 289-290
  - CREATE TABLE, 45-47, 50-51, 382
    - CUSTOMER TBL, 436
    - EMPLOYEE PAY TBL, 435
    - EMPLOYEE TBL, 435
    - ORDERS TBL, 436
    - PRODUCTS TBL, 436
  - CREATE TABLE AS, 382
  - CREATE TRIGGER, 349-350
  - CREATE TYPE, 382
  - CREATE VIEW, 382
    - GROUP BY clause, 323
    - views from multiple tables, 318-319
    - views from other views, 319-320
    - views from single tables, 316-318
    - WITH CHECK OPTION, 320-321
  - DELETE, 383
    - subqueries, 226
    - table data, 81
    - WHERE clause, 81
  - DISCONNECT, 14
  - DISTINCT, 142
  - DROP, 51
    - indexes, 260-261
    - users, 293
  - DROP INDEX, 383
  - DROP ROLE, 306
  - DROP SCHEMA, 290
  - DROP TABLE, 383
  - DROP TRIGGER, 351
  - DROP VIEW, 323, 383
  - EXIT, 14
  - FETCH, 345
  - formatting, 266
  - GRANT, 383
    - ADMIN OPTION, 303
    - GRANT OPTION, 303
    - privileges, 301
    - user access control, 302-303

## statements

- INSERT, 383
  - adding data to columns, 75-76
  - adding data to tables, 74-78
  - CUSTOMER TBL, 438
  - EMPLOYEE PAY TBL, 438
  - EMPLOYEE TBL, 437
  - ORDERS TBL, 439
  - PRODUCTS TBL, 440
  - subqueries, 224-225
- INSERT...SELECT, 383
- interactive, 375-376
- RELEASE SAVEPOINT, 94
- REVOKE, 384
  - privileges, 301
  - user access control, 303-304
  - users, 293
- ROLLBACK, 90-92, 384
- ROLLBACK TO SAVEPOINT, 92-94
- SAVEPOINT, 92, 384
- SELECT, 384
  - ALL option, 104
  - case sensitivity, 108
  - clauses, 102
  - column aliases, 112-113
  - COUNT function, 111
  - DISTINCT option, 104
  - EXCEPT operator, 241-242
  - FROM clause, 104
  - GROUP BY clause, 152-159, 244-245
  - HAVING clause, 159-160
  - implementation differences, 370
  - INTERSECT operator, 240-241
  - ORDER BY clause, 106-108, 242-244
  - queries, 101-104
  - selecting data from another table, 112
  - single queries, 235
  - subqueries, 223-224
  - UNION ALL operator, 239-240
  - UNION operator, 237-239
  - WHERE clause, 105-106
- SET ROLE, 307
- SET TRANSACTION, 94
- tuning, 265-266
  - formatting, 266
  - FROM clause table arrangement, 269
  - full table scans, 272
  - HAVING clause, 275
  - indexes, 275-276
  - join order, 269-270
  - large sort operations, 275
  - LIKE operator, 273
  - OR operator, 274-275
  - readability, 267-269
  - stored procedures, 275
  - tools, 276
  - WHERE clause condition, 270-271
  - wildcard placement, 273
- types
  - data administration, 17
  - data control, 16
  - defining database structures, 15
  - manipulating data, 16
  - selecting data, 16
  - transactional control, 17
- UPDATE, 384
  - multiple columns, 80
  - single columns, 79-80
  - subqueries, 225-226
  - table data, 79

**static SQL, 351****stored procedures**

- advantages, 348
- definition, 346
- MySQL syntax, 347-348
- nesting, 346
- Oracle syntax, 347-348
- overview, 347
- performance, 275
- SQL Server syntax, 347-348

**storing dates/times, 186**

- DATETIME elements, 186
- implementation-specific data types, 187
- standard data types, 186

**strings**

- characters
  - adding, 176-177
  - ASCII, 178
  - date conversions, 196-197
  - functions, 165
  - positions, 172
  - replacing, 169
- concatenation, 166-168
- conversions
  - character to numbers, 179-180
  - numeric to characters, 180-181
- fixed-length, 29
- lengths, 175
- literal, 33-34
- lowercases, 170
- NULL values, 175-176
- query searches, 174-175
- substrings, 166, 170-171
- translating, 166-169
- trimming, 173-174

- uppercase, 169
- varying-length, 29

**Structured Query Language. See SQL****subqueries**

- BETWEEN operator, 222
- correlated, 229-230
- definition, 221
- DELETE statement, 226
- embedded, 227-228
- INSERT statement, 224-225
- overview, 221-222
- rules, 222
- SELECT statement, 223-224
- syntax, 222
- UPDATE statement, 225-226

**SUBSTR function, 170-171****substrings, 166, 170-171****subtraction operator (-), 135****SUM function, 144-146****summarized data maintenance, 315-316****Sybase**

- Open Client/C Developers Kit, 363
- parameters, 376
- system catalog implementations, 334
- system privileges, 299
- users, creating, 288-289

**synonyms**

- creating, 324-325
- definition, 324
- dropping, 325
- names, 326
- overview, 324
- PRIVATE, 324
- PUBLIC, 324

**system catalog****system catalog**

- creating, 331
- data, 331-332
- definition, 329
- implementations, 333-334
- maintenance, 332
- overview, 330
- passwords, 338
- querying, 335-336
- table queries, 338
- updating, 337

**systems**

- analysts, 285
- client/server, 12
- date, 188
- privileges, 299-300
- web-based database, 12-13

**T****tables**

- aliases, 208
- arranging in FROM clauses, 269
- base, 214-215
- columns, 21, 44-45
  - adding, 48
  - adding data, 75-76
  - aliases, 112-113
  - attributes, editing, 48
  - auto-incrementing, adding, 48
  - averaging values, 146-147
  - cardinality, 260
  - check constraints, 55-56
  - counting values, 142-144

- dropping constraints, 56
- editing, 49
- foreign keys, 53-54
- index considerations, 258
- maximum values, 147
- minimum values, 147-148
- NOT NULL constraints, 55
- NULL values, 78-79
- ordering with numbers, 156
- primary keys, 52-53
- qualifying, 205
- totaling values, 144, 146
- unique constraints, 53
- updating, 79-80
- user access control, 304

**creating, 45-47**

- existing table, 50-51
- views, 322

**data**

- deleting, 81
- inserting, 74-75
  - inserting from another table, 76-78
  - inserting into specified columns, 75-76
  - inserting NULL values, 78-79
- populating, 74
- selecting from another table, 112
- updating, 79-80

**data examples in book, 18, 20****dropping, 51, 57****editing, 47-49****fields, 20****joins**

- base tables, 214-215
- Cartesian product, 215-217
- component locations, 204

- equijoins, 204-206
- multiple keys, 213-214
- natural, 206-207
- non-equijoins, 208-209
- outer, 210-211
- self, 212-213
- table aliases, 208
- types, 204
- names, 18, 47
- NULL values, 22
- parent/child relationships, 54
- primary keys, 21
- records, 21, 111
- relational databases, 11
- rows, 21, 45
  - averaging values, 146-147
  - counting, 142-144
  - maximum values, 147
  - minimum values, 147-148
  - totaling values, 144-146
- selecting data from multiple, 203
- system catalog, 338
- windowed table functions, 354
- TBL suffix, 18**
- TEXT data type, 30**
- third normal forms, 66**
- time zone function, 189**
- times**
  - adding to dates, 190-191
  - data types
    - implementation-specific, 187
    - standard, 186
  - date functions, 187
    - adding time, 190-191
    - comparing dates/times, 191

- current, 188
- miscellaneous, 192
- time zones, 189
- DATETIME elements, 186
- storing, 186
- tools**
  - front-end, 63
  - GUI, 293
  - performance, 276
  - web database access, 365
- Transact-SQL, 373**
- transactions**
  - controlling, 88-90
  - databases, 17
  - definition, 87
  - initiating, 94
  - overview, 87
  - savepoints
    - deleting, 94
    - names, 92
    - performance, 95
    - rolling back, 92-94
  - saving changes, 89-90
  - undoing, 90-92
- TRANSLATE function, 166-169**
- translating strings, 166-169**
- triggers**
  - creating, 349-350
  - definition, 349
  - dropping, 351
  - FOR EACH ROW syntax, 351
- trimming strings, 173-174**
- troubleshooting passwords, 308**



**tuning****tuning**

- databases, 266

- SQL statements

- definition, 265-266

- formatting, 266

- FROM clause table arrangement, 269

- full table scans, 272

- HAVING clause, 275

- indexes, 275-276

- join order, 269-270

- large sort operations, 275

- LIKE operator, 273

- OR operator, 274-275

- readability, 267-269

- stored procedures, 275

- tools, 276

- WHERE clause condition, 270-271

- wildcard placement, 273

**types**

- statements

- data administration, 17

- data control, 16

- defining database structures, 15

- manipulating data, 16

- selecting data, 16

- transactional control, 17

- data

- basic, 28

- BLOB, 30

- BOOLEAN, 34

- CHAR, 29

- date and time, 32-33, 186-187

- DECIMAL, 31-32

- definition, 27

- domains, 35

- DOUBLE PRECISION, 32

- fixed-length strings, 29

- FLOAT, 32

- floating-point decimals, 32

- integers, 32

- large objects, 30

- lengths, 37

- literal strings, 33-34

- NULL, 34

- numeric, 30-31

- REAL, 32

- TEXT, 30

- user-defined, 35

- VARCHAR, 29

- varying-length strings, 29

- indexes, 255

- composite, 257

- implicit, 257

- single-column, 256

- unique, 256-258

- joins

- equijoins, 204-206

- natural, 206-207

- non-equijoins, 208-209

- outer, 210-211

- self, 212-213

- users, 284

**U**

- undoing transactions, 90-92

- UNION ALL operator, 239-240

- UNION operator, 235-239

- unique column constraints, 53

**unique indexes, 256-258**

**UPDATE object privilege, 301**

**UPDATE statement, 384**

subqueries, 225-226

table data, 79-80

**UPDATE(column\_name) object privilege, 301**

**updating**

system catalog, 337

table data, 79-80

view data, 321

**UPPER function, 169**

**uppercase strings, 169**

**USAGE object privilege, 300**

**user-defined data types, 35**

**users**

access, controlling

columns, 304

GRANT statement, 302-303

groups of privileges, 305

PUBLIC database account, 304

REVOKE statement, 303-304

authIDs, 283

creating, 286-287

MySQL, 289

Oracle, 287-288

SQL Server, 288-289

Sybase, 288-289

data, system catalog, 332

deleting, 293

editing, 291

GUI tools, 293

logical database design considerations, 63

managing, 285, 298

roles/privileges, 285

schemas, 289-290

schemas, compared, 286

sessions, 292

types, 284

## V

**values**

ceiling and floor function, 178

exponential function, 178

NULL

adding to columns, 78-79

checking, 175-176

replacing, 176

tables, 22

**VARCHAR data type, 29**

**varying-length strings, 29**

**vendors**

databases, 13-14

implementations, 369-371

**views**

creating, 316

multiple tables, 318-319

other views, 319-320

single tables, 316-318

WITH CHECK OPTION, 320-321

creating tables from, 322

data updates, 321

definition, 313

dependencies, 320

dropped tables, 326

dropping, 323

ORDER BY clause, 323

476

## views

- overview, 314
- rows, 321
- security, 315
- simplifying data, 314
- summarized data maintenance, 315-316

## W

**web interfaces, 363-364**

**web-based database systems, 12-13**

### websites

- ASCII chart, 178
- book, 9
- Center for Internet Security, 298
- MySQL, 375

**WHERE clause, 385**

- DELETE statement, 81
- restrictive condition, 270-271
- SELECT statement, 105-106

**wildcard performance, 273**

**windowed table functions, 354**

**WITH CHECK OPTION (CREATE VIEW statement),  
320-321**

## X-Z

**XML, 354-355**